

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Applicant:	Ambuj Shatdal	§	Art Unit:	2162
		§		
Serial No.:	10/751,017	§	Confirmation No.:	3663
		§		
Filed:	December 31, 2003	§	Examiner:	Jean B. Fleurantin
		§		
For:	Computing Aggregates on Distinct Attribute Values	§	Atty. Dkt. No.:	11087 (NCR.0119US)
		§		

Mail Stop Appeal Brief-Patents

Commissioner for Patents

P.O. Box 1450

Alexandria, VA 22313-1450

APPEAL BRIEF PURSUANT TO 37 C.F.R § 41.37

Sir:

The final rejection of claims 1-3, 8-11 and 15-20 is hereby appealed.

I. REAL PARTY IN INTEREST

The real party in interest is Teradata Corporation, previously affiliated with NCR Corporation.

II. RELATED APPEALS AND INTERFERENCES

None.

III. STATUS OF THE CLAIMS

Claims 1-3, 8-11 and 15-20 have been finally rejected and are the subject of this appeal.

Claims 4-7 and 12-14 were indicated as containing allowable subject matter.¹

¹ Since claims 4, 5, and 12 are independent claims, and claims 6, 7, 13 and 14 depend from respective base claims 5 and 12, it is believed that the Examiner intended to indicate these claims as allowed.

IV. STATUS OF AMENDMENTS

No Amendment was filed after the final rejection dated April 2, 2009—therefore, all amendments have been entered.

V. SUMMARY OF THE CLAIMED SUBJECT MATTER

The following provides a concise explanation of the subject matter defined in each of the independent claims involved in the appeal, referring to the specification by page and line number and to the drawings by reference characters, as required by 37 C.F.R. § 41.37(c)(1)(v). Each element of the claims is identified by a corresponding reference to the specification and drawings where applicable. Note that the citation to passages in the specification and drawings for each claim element does not imply that the limitations from the specification and drawings should be read into the corresponding claim element.

Independent claim 1 is directed to computing aggregates over distinct attribute values in response to a query that specifies an aggregate on distinct values of at least one attribute, where the query further specifies grouping on plural grouping sets including a first grouping set and a second grouping set. Specification, p. 6, ¶ [0023], ln. 1-7; ¶ [0024], ln. 1-4. Distinct values of the at least one attribute are identified, and such distinct values are stored in a first table. Specification, p. 7, ¶ [0025], ln. 1-5; ¶ [0026], ln. 1-6. Aggregates for groups specified by the first grouping set and the second grouping set are computed using the first table that contains the distinct values of the at least one attribute. Specification, p. 10, ¶ [0036], ln. 1; ¶ [0038], ln. 9. Specifically, claim 1 recites a method executable by at least one processor in a database system, comprising:

receiving, by the at least one processor, a query that specifies an aggregate on distinct values of at least one attribute, the query further specifying grouping on plural grouping sets, the plural grouping sets having at least a first grouping set and a second grouping set (Fig. 1:110; Spec., ¶¶ [0023], [0024], [0027], [0028], [0031]);

identifying, by the at least one processor, distinct values of the at least one attribute and storing the distinct values of the at least one attribute in a first table (Fig. 3:206; Spec., ¶¶ [0025], [0026], [0031], [0035]);

computing, by the at least one processor, aggregates for groups specified by the first grouping set using the first table (Fig. 3:208; Spec., ¶¶ [0036], [0037], [0044]); and

computing, by the at least one processor, aggregates for groups specified by the second grouping set using the first table (Fig. 3:210; Spec. ¶¶ [0038], [0041]-[0044]).

Independent claim 8 is directed to computing aggregates over distinct attribute values in response to a query that specifies an aggregate on distinct values of at least one attribute, where the query further specifies grouping on plural grouping sets including a first grouping set and a second grouping set. Specification, p. 6, ¶ [0023], ln. 1-7; ¶ [0024], ln. 1-4. Distinct values of the at least one attribute are identified, and such distinct values are stored in a first table. Specification, p. 7, ¶ [0025], ln. 1-5; ¶ [0026], ln. 1-6. Aggregates for groups specified by the first grouping set and the second grouping set are computed using the first table that contains the distinct values of the at least one attribute. Specification, p. 10, ¶ [0036], ln. 1; ¶ [0038], ln. 9. Specifically, claim 8 recites an article comprising at least one machine-readable storage medium containing instructions that when executed cause a system to:

receive a query that specifies an aggregate on distinct values of at least one attribute, the query further specifying grouping on plural grouping sets, the plural grouping sets having at least a first grouping set and a second grouping set (Fig. 1:110; Spec., ¶¶ [0023], [0024], [0027], [0028], [0031]);

identify distinct values of the at least one attribute and storing the distinct values of the at least one attribute in a first table (Fig. 3:206; Spec., ¶¶ [0025], [0026], [0031], [0035]);

compute aggregates for groups specified by the first grouping set using the first table (Fig. 3:208; Spec., ¶¶ [0036], [0037], [0044]); and

compute aggregates for groups specified by the second grouping set using the first table (Fig. 3:210; Spec. ¶¶ [0038], [0041]-[0044]).

Independent claim 15 is directed to a database system that computes aggregates over distinct attribute values in a table in response to a query that specifies group-by operations on plural grouping sets. Specification, p. 6, ¶ [0023], ln. 1-7; ¶ [0024], ln. 1-4. In processing the query, intermediate values are computed for storage in an intermediate spool, and the intermediate values are used for computing results of at least two group-by operations on at least two corresponding grouping sets. Specification, p. 7, ¶ [0025], ln. 1-5; ¶ [0026], ln. 1-6; p. 10, ¶ [0036], ln. 1; ¶ [0038], ln. 9. Specifically, claim 15 recites a database system (Fig. 1:10) comprising:

a machine-readable storage (Fig. 1:104) to store a table; and

at least one processor (Fig. 1:112) to:

receive a query that specifies a calculation of an aggregate on distinct values of an attribute in the table, the query to specify group-by operations on plural grouping sets (Fig. 1:110; Spec., ¶¶ [0023], [0024], [0027], [0028], [0031]);

in processing the query, compute intermediate values for storage in an intermediate spool (Fig. 3:206; Spec., ¶¶ [0025], [0026], [0031], [0035]); and

use the intermediate values in the intermediate spool for computing results of at least two group-by operations on at least two corresponding grouping sets (Fig. 3:208, 210; Spec., ¶¶ [0036]-[0038], [0041]-[0046]).

VI. GROUNDS OF REJECTION TO BE REVIEWED ON APPEAL

- A. Claims 1-3, 8-11 and 15-20 were rejected under 35 U.S.C. § 103(a) as unpatentable over U.S. Patent No. 5,713,020 (Reiter) in view of Applicant background, specification, pages 2-4, up to ¶ [0009] (APA).**

VII. ARGUMENT

The claims do not stand or fall together. Instead, Appellant presents separate arguments for various independent and dependent claims. Each of these arguments is separately argued below and presented with separate headings and sub-headings as required by 37 C.F.R. § 41.37(c)(1)(vii).

- A. Claims 1-3, 8-11 and 15-20 were rejected under 35 U.S.C. § 103(a) as unpatentable over U.S. Patent No. 5,713,020 (Reiter) in view of Applicant background, specification, pages 2-4, up to [0009] (APA).**

1. Claims 1, 2, 8, 9.

It is respectfully submitted that the obviousness rejection of claim 1 over Reiter and APA is erroneous.

To make a determination under 35 U.S.C. § 103, several basic factual inquiries must be performed, including determining the scope and content of the prior art, and ascertaining the differences between the prior art and the claims at issue. *Graham v. John Deere Co.*, 383 U.S. 1, 17, 148 U.S.P.Q. 459 (1965). Moreover, as held by the U.S. Supreme Court, it is important to identify a reason that would have prompted a person of ordinary skill in the art to combine reference teachings in the manner that the claimed invention does. *KSR International Co. v. Teleflex, Inc.*, 127 S. Ct. 1727, 1741, 82 U.S.P.Q.2d 1385 (2007).

The following describes the teachings of the cited prior art, Reiter and APA, and how they differ from the claimed subject matter. Claim 1 recites a method executable by at least one processor in a database system that comprises:

- receiving, by the at least one processor, a query that specifies an aggregate on *distinct values* of at least one attribute, the query further specifying grouping on plural grouping sets, the plural grouping sets having at least a first grouping set and a second grouping set;
- identifying, by the at least one processor, *distinct values* of the at least one attribute and storing the *distinct values* of the at least one attribute in a first table;
- computing, by the at least one processor, aggregates for groups specified by the first grouping set using the first table; and
- computing, by the at least one processor, aggregates for groups specified by the second grouping set using the first table.

The Examiner cited column 4, lines 34-37, of Reiter as disclosing a query that specifies an aggregate on *distinct values* of at least one attribute. 4/2/2009 Office Action at 5. Column 4 of Reiter refers to a multiple-level aggregation query, such as in the form of Table 1, which includes a GROUP BY clause that specifies a level number. Column 5 of Reiter shows an example query (*see* Table 2 in column 5 of Reiter). In this example query, several aggregations are specified, all of which pertain to the SUM() aggregation. The aggregation is performed on the Order Amount attribute, with absolutely *no* indication that the SUM is performed on *distinct* values of the Order Amount attribute. Therefore, contrary to the assertion by the Examiner, it is clear that Reiter does not teach or hint at receiving a query that specifies an aggregate on distinct values of at least one attribute.

In fact, the query in column 5 of Reiter sums *all* attribute values, not just *distinct* values. For example, as stated in column 5, at lines 24-25, of Reiter: “the query contains a sum of order amounts for *all* of the rows in the source table.” In column 2, in connection with the discussion

of an aggregate function that can be contained in a SQL query, Reiter states that the aggregating function “performs some operation on the values of *all* of the fields in that column from the rows of the source table” Reiter, 2:24-27. Thus, it is clear that the queries referred to in Reiter are queries that aggregate all values of at least one attribute, not *distinct* values of at least one attribute, to form groups.

The Examiner also made the following observation with respect to Reiter:

Moreover, Reiter discloses a query is a set of instructions to combine, filter, and sort the contents of one or more tables. SQL permits the specification of a query with a single level of aggregation. Such a query specifies grouping columns. For every unique combination of fields in the grouping column of a source table, a single row is created in the query table. For each such field, an aggregating function is designated. The aggregating function is used to produce a value for each field in the aggregated column.

4/2/2009 Office Action at 3.

This statement regarding Reiter does not provide any explanation that Reiter discloses receiving a query that specifies an aggregate on **distinct** values of at least one attribute. The fundamental difference between the teachings of Reiter and the subject matter recited in claim 1, strongly indicates that the hypothetical combination of Reiter and APA clearly does not teach or hint at elements of claim 1.

In fact, there clearly did not exist any reason to combine the teachings of Reiter and APA to achieve the claimed invention. As noted by a recent U.S. Supreme Court case, it is important to identify a reason that would have prompted a person of ordinary skill in the art to combine the teachings of the cited references to achieve the claimed invention. *KSR*, 127 S. Ct. at 1741. In fact, as objective evidence that no reason existed to combine Reiter and APA, the APA actually teaches away from the invention. The Examiner conceded that Reiter fails to disclose storing distinct values of the at least one attribute (specified in the query) in a first table. 4/2/2009 Office Action at 5. Instead, the Examiner relied upon APA, and specifically, to page 3, ¶ [0006],

line 1 of the APA (the Background section of the present application). The APA, in ¶ [0007], provides an explicit teaching that the approach described in ¶¶ [0005] and [0006] of APA “cannot be used if the SQL statement specifies that an aggregate be calculated on distinct values of a particular attribute” Thus, the APA expressly taught away from using the technique in ¶ [0006] of APA in the context of what is recited in claim 1, namely in the context of a query that specifies an aggregate on distinct values of at least one attribute. Thus, the APA provides objective evidence that there existed no reason for a person of ordinary skill in the art to combine the teachings of Reiter and APA in the manner proposed by the Examiner.

In response to Appellant’s arguments, the Examiner argued that claims are to be given “the broadest reasonable interpretation consistent with the specification” 4/2/2009 Office Action at 4. This statement does not address the issue of Appellant’s argument that no reason existed to combine Reiter and APA in view of the teaching away by APA. As stated by the M.P.E.P., “[a] prior art reference must be considered in its entirety, *i.e.*, as a whole, including portions that would lead away from the claimed invention.” M.P.E.P. § 2141.02 (8th ed., Rev. 6), at 2100-126. In fact, “[i]t is improper to combine references where the references teach away from their combination.” M.P.E.P. § 2145, at 2100-167-168; *see also United States v. Adams*, 383 U.S. 39, 51-52 (1966) (holding that prior art would not render obvious claimed subject matter if the prior art taught away from the claimed invention). Thus, in view of the foregoing, there existed no reason for a person of ordinary skill in the art to combine the teachings of Reiter and APA in the manner proposed by the Examiner.

Therefore, it is respectfully submitted that independent claim 1 and its dependent claims are non-obvious over Reiter and APA.

Independent claim 8 and its dependent claims are non-obvious over Reiter and APA for similar reasons as claim 1.

Reversal of the final rejection of the above claims is respectfully requested.

2. Claims 3, 10, 11.

Claims 3, 10, and 11 depend from claims 1 and 8, respectively, and are allowable for at least the same reasons as claim 1. Moreover, claim 3 recites that identifying the distinct values of the at least one attribute comprises computing a group-by operation on the first grouping set and selecting the attributes of the first grouping set for output. The Examiner cited Table 2 in column 5 of Reiter, and specifically, lines 6 and 7 of Table 2, as disclosing this feature of claim 3. The SQL query of Table 2 clearly does not disclose a query to identify *distinct values* by computing a group-by operation in the manner recited in claim 3. In fact, the SELECT clauses at lines 2 and 3 of Table 2, which correspond to the group-by clauses at lines 6 and 7 of Table 2, perform the SUM aggregation on the Order Amount attribute. Therefore, the result of the operation of Table 2 is that an aggregation (sum) is provided, not distinct values of the attribute, as recited in claim 3.

For the foregoing reasons, claim 3 is further allowable over the cited references.

Claims 10 and its dependent claim 11 are allowable for similar reasons as claim 3.

Reversal of the final rejection of the above claims is respectfully requested.

3. Claims 15-20.

Independent claim 15 recites a database system having a machine-readable storage to store a table and at least one processor to:

- receive a query that specifies a calculation of an aggregate on distinct values of an attribute in the table, the query to specify group-by operations on plural grouping sets;
- in processing the query, compute intermediate values for storage in an intermediate spool; and
- use the intermediate values in the intermediate spool for computing results of at least two group-by operations on at least two corresponding grouping sets.

As discussed above in connection with claim 1, the Examiner has mis-applied Reiter against the “receive” element, which recites “receive a query that specifies a calculation of an aggregate on *distinct values* of an attribute in the table, the query to specify group-by operations on plural grouping sets.” Moreover, as explained above in connection with claim 1, there existed no reason to combine the teachings of Reiter and APA, and in fact, the APA teaches away from the combination.

In view of the foregoing, it is clear that a *prima facie* case of obviousness has also not been established with respect to claim 15 and its dependent claims.

Reversal of the final rejection of the above claims is respectfully requested.

CONCLUSION

In view of the foregoing, reversal of all final rejections and allowance of all pending claims is respectfully requested.

Respectfully submitted,

Date: September 2, 2009

/Dan C. Hu/

Dan C. Hu
Registration No. 40,025
TROP, PRUNER & HU, P.C.
1616 South Voss Road, Suite 750
Houston, TX 77057-2631
Telephone: (713) 468-8880
Facsimile: (713) 468-8883

VIII. APPENDIX OF APPEALED CLAIMS

The claims on appeal are (claims 4-7 and 12-14 are indicated as containing allowable subject matter):

- 1 1. A method executable by at least one processor in a database system, comprising:
2 receiving, by the at least one processor, a query that specifies an aggregate on distinct
3 values of at least one attribute, the query further specifying grouping on plural grouping sets, the
4 plural grouping sets having at least a first grouping set and a second grouping set;
5 identifying, by the at least one processor, distinct values of the at least one attribute and
6 storing the distinct values of the at least one attribute in a first table;
7 computing, by the at least one processor, aggregates for groups specified by the first
8 grouping set using the first table; and
9 computing, by the at least one processor, aggregates for groups specified by the second
10 grouping set using the first table.
- 1 2. The method of claim 1, wherein the first grouping set is lower level grouping set than the
2 second grouping set, and wherein the first grouping set has a larger number of attributes than the
3 second grouping set.
- 1 3. The method of claim 1, wherein identifying the distinct values of the at least one attribute
2 comprises computing a group-by operation on the first grouping set and selecting the attributes
3 of the first grouping set for output.

1 8. An article comprising at least one machine-readable storage medium containing
2 instructions that when executed cause a system to:
3 receive a query that specifies an aggregate on distinct values of at least one attribute, the
4 query further specifying grouping on plural grouping sets, the plural grouping sets having at least
5 a first grouping set and a second grouping set;
6 identify distinct values of the at least one attribute and storing the distinct values of the at
7 least one attribute in a first table;
8 compute aggregates for groups specified by the first grouping set using the first table; and
9 compute aggregates for groups specified by the second grouping set using the first table.

1 9. The article of claim 8, wherein the first grouping set is lower level grouping set than the
2 second grouping set, and wherein the first grouping set has a larger number of attributes than the
3 second grouping set.

1 10. The article of claim 8, wherein identifying the distinct values of the at least one attribute
2 comprises computing a group-by operation on the first grouping set and selecting the attributes
3 of the first grouping set for output.

1 11. The article of claim 10, wherein storing the distinct values of the at least one attribute in
2 the first table comprises storing the distinct values of the at least one attribute in a spool file.

1 15. A database system comprising:
2 a machine-readable storage to store a table; and
3 at least one processor to:
4 receive a query that specifies a calculation of an aggregate on distinct values of an
5 attribute in the table, the query to specify group-by operations on plural grouping sets;
6 in processing the query, compute intermediate values for storage in an
7 intermediate spool; and
8 use the intermediate values in the intermediate spool for computing results of at
9 least two group-by operations on at least two corresponding grouping sets.

- 1 16. The database system of claim 15, wherein the query comprises a Structured Query
2 Language (SQL) SELECT statement containing a GROUP BY clause specifying multiple
3 grouping sets.
- 1 17. The database system of claim 15, wherein the query specifies group-by operations on
2 plural grouping sets at multiple grouping levels.
- 1 18. The database system of claim 15, further comprising database management software
2 executable on the at least one processor to perform the receiving, computing, and using acts.
- 1 19. The database system of claim 18, wherein the database management software comprises
2 plural access modules, and the storage comprises plural storage modules accessible by the plural
3 access modules in parallel.
- 1 20. The database system of claim 19, further comprising plural processors, the access
2 modules executable on the processors.

IX. EVIDENCE APPENDIX

None.

X. RELATED PROCEEDINGS APPENDIX

None.